

# Fast Sun-aligned Outdoor Scene Relighting based on TensoRF

Yeonjin Chang<sup>1</sup> Yearim Kim<sup>1</sup> Seunghyeon Seo<sup>1</sup> Jung Yi<sup>2</sup> Nojun Kwak<sup>1</sup>

<sup>1</sup>Seoul National University

{yjean8315, yerim1656, zzzlssh, nojunk}@snu.ac.kr

<sup>2</sup>Yonsei University

master9906@yonsei.ac.kr

## Abstract

In this work, we introduce our method of outdoor scene relighting for Neural Radiance Fields (NeRF) named Sun-aligned Relighting TensoRF (SR-TensoRF). SR-TensoRF offers a lightweight and rapid pipeline aligned with the sun, thereby achieving a simplified workflow that eliminates the need for environment maps. Our sun-alignment strategy is motivated by the insight that shadows, unlike viewpoint-dependent albedo, are determined by light direction. We directly use the sun direction as an input during shadow generation, simplifying the requirements of the inference process significantly. Moreover, SR-TensoRF leverages the training efficiency of TensoRF by incorporating our proposed cubemap concept, resulting in notable acceleration in both training and rendering processes compared to existing methods.

## 1. Introduction

Recently, Neural Radiance Fields (NeRF) [26] has achieved a breakthrough in the novel view synthesis task, sparking significant interest within the academic community in implicit 3D reconstruction. While NeRF’s impressive performance is undeniable, numerous challenges remain for further enhancement and research, such as reconstructing materials with various properties like transparency, translucency, or reflectiveness [12, 13, 38, 43, 47], reconstructing dynamic scenes [1, 21, 30, 40] or scenes with varying scales [3, 37, 42], efficient training and inference [8, 9, 11, 27, 49], and more. In addition to the aforementioned challenges, another area of interest lies in the editing of reconstructed scenes [15, 52]. One aspect of such editing involves manipulating lighting conditions [32, 36, 39], which is referred to as relighting.

Relighting is a classical task in computer vision that enhances the realism and versatility of visual content, benefiting applications like virtual environments, entertainment media, and architectural visualization. Due to the intrinsic difference in lighting between indoor and outdoor scenes,

researchers have treated them as separate domains and categorized the relighting problem into indoor and outdoor scene relighting. In this study, we focus on the daytime outdoor setting which can be characterized as; 1) Sunlight stands as the most dominant light source. 2) Outdoor settings lack the ability to control light sources contrary to indoor settings utilizing controllable light conditions. 3) Outdoor scenes are inherently unbounded so that they are profoundly influenced by the surrounding objects and buildings while consistently dealing with infinite potential light sources. 4) Training images might have varying white balances, even within the same timeframe.

Recently, with the purpose of achieving the dual objectives of novel view synthesis and relighting, NeRF-

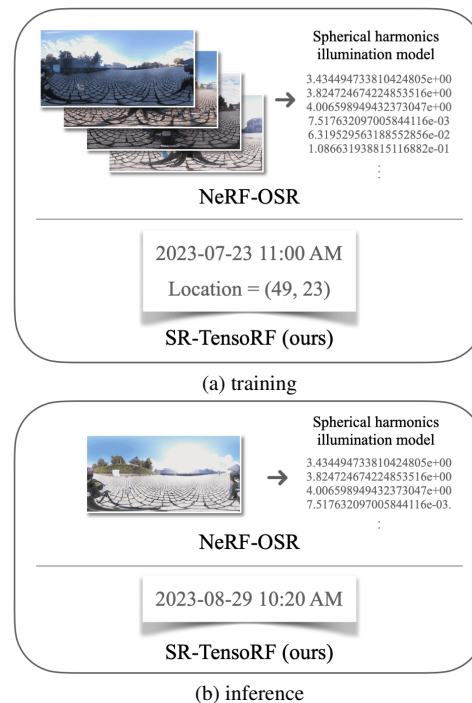


Figure 1. **Our requirements.** Instead of using environment maps, we propose a simplified relighting pipeline that relies on time and location information, making the requirements much simpler.

OSR [32], which decomposes the implicit scene using shadows, albedo, and an illumination model based on spherical harmonics, has been proposed. While the approach generated plausible outputs with their relighting dataset, confirming its alignment with the aforementioned outdoor scene features is challenging. Moreover, during the inference process, it requires the transformation of environment maps to spherical harmonics. To attain a specific lighting condition, it might be necessary to capture an environment map tailored exclusively to that particular condition. While NeRF is progressing towards becoming a more lightweight and faster model, employing environment maps hinders the efficiency of the combined task of NeRF and relighting. Hence, particularly for daytime outdoor scenes, we present the feasibility of relighting without environment maps.

In this paper, we leverage the fact that sunlight is the most dominant light source in outdoor scenes. Therefore, an effective approach to relighting in outdoor settings involves accurately generating shadows cast by the sun. In our pursuit of simultaneous novel view synthesis and relighting, we organize the synthesis process into three core components: *albedo (color)*, *shadow*, and *tone (brightness)*. The acquisition of albedo closely resembles the learning process for color in existing NeRF methods. Shadows, governed by the sun direction and the surrounding geometry, are learned by taking the sun direction and shadow features as input, allowing the training of time-dependent shadows. While albedo and shadow already form the foundational elements of relighting, it is important to note that training images may exhibit varying tones even within the same time frame. To disentangle the influence of tone on the learning of albedo and shadows, we introduce a latent vector that deals with tone throughout the training process.

Another key aspect of our model lies in its foundation architected atop TensorRF [8], which is known for its efficient training capabilities. However, since TensorRF was not originally intended for relighting, its performance in relighting datasets falls short of satisfaction. To address this limitation, we have devised an architecture based on TensorRF to the specific demands of the relighting task to achieve fast and effective relighting. This tailored approach endows both our training and rendering procedures an appreciable acceleration, surpassing NeRF-OSR [32] by a substantial margin.

We summarize our contributions as follows:

- **No need of environment map** Our model performs relighting based primarily on the sun direction, thereby eliminating the complexity of using an environment map in training and inference processes (see Fig. 1).
- **Handle unbounded scene** Our model employs cube-map based on TensorRF to successfully model unbounded scenes.
- **Fast relighting and rendering** Our model achieves a speedup of around tenfold in training and inference processes compared to previous research.

## 2. Related Works

### 2.1. Neural Radiance Field

Neural Radiance Field (NeRF) [26] has been mainstream for the novel view synthesis task thanks to its simple architecture and splendid quality of rendered images. NeRF utilizes a Multi-Layer Perceptron (MLP) to depict a scene by associating 3D coordinates and viewing angles with colors and volume densities, and generates a novel view using the volume rendering technique. Since its introduction, NeRF has been actively studied toward various directions, *e.g.*, a multi-scale representation [2, 4, 17], data efficiency [28, 33, 34, 50], dynamic scene reconstruction [1, 21, 30, 40], and so on. Furthermore, there has been a line of research to address the challenging problems resulting from the fundamental limitations of NeRF’s architecture, *e.g.*, lengthy training time per scene [8, 9, 11, 27, 49], incapacity to handle unbounded scenes [3, 37, 42], novel illumination conditions [6, 22, 32, 36, 39], and so on. Among them, TensorRF [8] represents a scene with a voxel feature grid decomposed into several low-rank tensor components, achieving a more compact scene representation. In this paper, we address the representative bottlenecks altogether and propose SR-TensorRF, which performs relighting successfully for unbounded outdoor scenes with a significantly reduced training time compared to other baselines.

### 2.2. Editing Neural Radiance Field

While NeRF achieves notable performance in generating photo-realistic novel views, editing them remains quite challenging since the 3D scene is represented by MLP parameters in its framework, which are hard to interpret and manipulate. There exist various lines of research with regard to editing NeRF, *e.g.*, scene stylization [15, 16, 45, 53], object-level editing [19, 35, 48], and so on. Among them, a series of studies [18, 24, 36, 39] have addressed the task of controlling attributes of objects or scenes, *e.g.*, illumination, material, and so on. TensorIR [18] expanded upon TensorRF [8] to enable material editing and relighting with high-quality rendering. However, it assumed a set of relatively simple scenes with controlled lighting or simple objects, which are far from real-world scenarios. Our proposed SR-TensorRF empowers controlling on more complex unbounded outdoor scenes, all while leveraging the effective representation of TensorRF.

### 2.3. Outdoor Scene Relighting

There have been various studies on outdoor scene relighting, emphasizing its importance in graphics and VR/AR. Traditional researches [10, 14, 41] optimized and computed variables without training. Duchêne *et al.* [10] used a 3D reconstruction with sun direction to segregate images into reflectance and shading, requiring manual shadow




Figure 2. **Overview.** SR-TensoRF comprises three tensors:  $\mathcal{G}_\sigma$ ,  $\mathcal{G}_\alpha$ , and  $\mathcal{G}_c$ , among which  $\mathcal{G}_\sigma$  learns features of density,  $\mathcal{G}_\alpha$  is for both albedo and shadow, and  $\mathcal{G}_c$  is for cubemap, respectively. These components, along with the viewing direction  $\mathbf{d}$ , sun direction  $\mathbf{d}_{sun}$ , and latent vectors, are processed through shallow MLPs to generate albedo, shadow, and cubemap backgrounds. These generated components are then used in volumetric rendering to produce the final output, where the  $\odot$  operator denotes pixel-wise multiplication. The green ray depicts the standard volumetric rendering, while the blue ray represents a case where the accumulated transmittance of sampled points on the ray is low (small  $\sigma$  over the entire ray-marching), necessitating the cubemap.

clicking to determine sun direction. After the advent of deep learning, researches [7,23,29,51] leveraged neural networks to achieve relighting. Yu *et al.* [51] used a self-supervised method with a neural renderer and a dedicated network for shadow prediction, but it lacked accurate color representation due to the absence of color correction.

After NeRF’s introduction, NeRF-OSR [32] extended a novel view synthesis framework to handle outdoor scene relighting, which augmented NeRF’s pipeline for learning shadow and normal decomposition. However, it does not consider the characteristics of outdoor scenes and requires a lengthy training time. Unlike the previous methods, we enable a much more efficient pipeline by using solar information in the training process, skipping the pre-processing step of creating an environment map.

### 3. Method

In this section, we present our approach, SR-TensoRF, to efficiently achieve simultaneous relighting and novel view synthesis in outdoor scenes, as shown in Figure 2. SR-TensoRF is designed based on TensoRF [8], which employs 4D tensor decomposition to model voxel grids (Sec. 3.1). Instead of environment maps, we leverage sun direction directly, accounting for its variation across time and location, to facilitate relighting (Sec. 3.2). For unbounded scenes not encompassed within TensoRF’s architecture, we incorporate cubemaps (Sec. 3.3). Furthermore, we utilize latent vectors to disentangle the non-uniform color tones present in the training images (Sec. 3.4). We optimize SR-TensoRF using reconstruction loss, TensoRF’s training loss, and shadow regularization loss (Sec. 3.5).

### 3.1. Preliminary

#### 3.1.1 Neural Radiance Field

NeRF [26] employs a neural network to depict a 3D scene with the 3D coordinates of a point  $\mathbf{x}$  and the viewing direction  $\mathbf{d}$  as input, producing the color  $\mathbf{c}$  and density  $\sigma$  of the points as output. The input 3D points are sampled along a ray  $\mathbf{r}(t) = \mathbf{o} + t\mathbf{d}$ , and each ray is cast toward a pixel in training images. The outputs, namely the color  $\mathbf{c}$  and density  $\sigma$ , are then utilized to render a pixel color  $\hat{\mathbf{C}}(\mathbf{r})$ :

$$\hat{\mathbf{C}}(\mathbf{r}) = \sum_{i=1}^N T_i (1 - \exp(-\sigma_i \delta_i)) \mathbf{c}_i = \sum_{i=1}^N w_i \mathbf{c}_i \quad (1)$$

$$w_i = T_i (1 - \exp(-\sigma_i \delta_i)), \quad T_i = \exp\left(-\sum_{j=1}^{i-1} \sigma_j \delta_j\right)$$

where  $N$  is the number of sample points, and  $\delta_i$  and  $T_i$  are the step size of a ray and the accumulated transmittance at each sampled point  $\mathbf{r}(t_i)$ , respectively.

#### 3.1.2 Tensorial Radiance Fields (TensoRF)

TensoRF [8] represents a 3D scene with a 3D geometry tensor  $\mathcal{G}_\sigma$  and a 4D appearance tensor  $\mathcal{G}_\alpha$ , where an additional dimension corresponds to the 3D point’s 1D appearance feature. The element at each 3D position  $\mathbf{x}$  of  $\mathcal{G}_\sigma$  is a density  $\sigma(\mathbf{x})$  while that of  $\mathcal{G}_\alpha$  is a 1D feature vector. This feature vector, along with the viewing direction  $\mathbf{d}$ , is fed into a shallow MLP or a spherical harmonics function  $M$ , ultimately generating the color  $\mathbf{c}$ :

$$\sigma, \mathbf{c} = \mathcal{G}_\sigma(\mathbf{x}), M(\mathcal{G}_\alpha(\mathbf{x}), \mathbf{d}). \quad (2)$$




Figure 3. **Calculating the sun direction from time and location information.** Using PyEphem [31], we get the sun directions from the local observer’s coordinates and subsequently align them with the first camera’s coordinates.

The key advantage of TensorRF lies in its lightweight model architecture and rapid training speeds. This is achieved by employing an alpha mask for densities below a specific threshold, ensuring that only the minimum required number of remaining voxels is retained. The voxel-wise alpha value is computed as  $\alpha = 1 - \exp(-\sigma\delta)$ . For more architectural details, kindly refer to the paper of TensorRF.

### 3.2. Sun Light Direction as Input

For successful daytime outdoor relighting, it is essential to generate shadows caused by sunlight accurately. These shadows are primarily determined by the sun direction and are influenced by the surrounding geometry while remaining independent of the viewing direction. Taking these aspects into consideration, we organize the process of generating shadows in two steps: the first involves learning *geometry* using a shadow feature from a 4D appearance tensor, and the second is using the *sun direction as input* of the shading process.

We implement the first step by making the 4D appearance tensor  $\mathcal{G}_a$  incorporate not only the albedo but also the shadow, *i.e.*,  $\mathcal{G}_a(\mathbf{x}) = \mathcal{G}_{albedo}(\mathbf{x}) \oplus \mathcal{G}_{shadow}(\mathbf{x})$ . The shadow tensor  $\mathcal{G}_{shadow}$  is intended to depict intricate shadowing conditions influenced by the surrounding geometry. Later, in the second step, the shadow feature,  $\mathcal{G}_{shadow}(\mathbf{x})$ , is input into the shadow MLP  $M_{shadow}$  alongside the sun

direction  $\mathbf{d}_{sun}$  as follows:

$$\begin{aligned} \sigma &= \mathcal{G}_\sigma(\mathbf{x}) \\ \mathbf{c} &= M_{albedo}(\mathcal{G}_{albedo}(\mathbf{x}), \mathbf{d}) \\ s &= M_{shadow}(\mathcal{G}_{shadow}(\mathbf{x}), \mathbf{d}_{sun}) \end{aligned} \quad (3)$$

where  $s \geq 0$  indicates the strength of shadow at 3D sampled point  $\mathbf{x}$ . The  $\sigma$ ,  $\mathbf{c}$ , and  $s$  obtained from Eq. (3) are used to volume render albedo  $\mathbf{A}(\mathbf{r})$  with  $(\sigma, \mathbf{c})$  and similarly, to render shadow  $S(\mathbf{r}) \in [0, 1]$  with  $(\sigma, s)$  as in Eq. (1). Subsequently, the two results are multiplied pixel-wise to generate the final output  $\hat{\mathbf{C}}(\mathbf{r})$ .

The sun direction  $\mathbf{d}_{sun}$  can be computed from the date, time, and location coordinates of the captured images using PyEphem [31] as in Fig. 3. Once determined, this sun direction, being in the horizontal coordinates, which uses the observer’s local horizon as the x-z plane, must be aligned with the world coordinates of the training images. Given the convention of the NeRF camera pose where the camera pose of the first training image serves as the world coordinate, the sun direction is rotated along the y-axis based on how much the first training image deviates from the south.

### 3.3. Cubemap to Handle Unbounded Scenes

TensorRF encounters two types of challenges when reconstructing unbounded scenes due to its limited ability to effectively utilize alpha masks. When alpha masks do not mask out densities sufficiently, it can lead to an elongated training time, the unintentional learning of extraneous densities, and, consequently, the emergence of artifacts. On the other hand, when alpha masks function excessively, it can result in cracks in the necessary background regions of unbounded scenes. Hence, we incorporate the concept of cubemaps in computer graphics to handle backgrounds into TensorRF. Cubemaps are beneficial at capturing missing background information, resembling the role of background networks, especially as seen in prior NeRF research like NeRF++ [54]. However, since TensorRF parametrizes scenes using cuboid shapes, cubemaps offer a more seamless and suitable background.

To integrate our proposed cubemap, we add background to the volume rendering equation (1):

$$\begin{aligned} \hat{\mathbf{C}}(\mathbf{r}) &= \sum_{i=1}^N w_i \mathbf{c}_i + (1 - w_i) \mathbf{b} \\ \mathbf{b} &= M_{cubemap}(\mathcal{G}_c(\mathbf{r}), \mathbf{d}_{sun}), \end{aligned} \quad (4)$$

where  $\mathbf{b}$  is the cubemap background,  $M_{cubemap}$  is a shallow MLP, and  $\mathcal{G}_c(\mathbf{r})$  is the cubemap feature bilinearly interpolated from where the ray intersects with one of the six 3D cubemap tensors  $\mathcal{G}_c$ .  $M_{cubemap}$  is conditioned by the sun direction since background colors might vary based on it. Furthermore, to address the inherent TensorRF’s inability to

model unbounded scenes, we have integrated the ray contraction technique from MipNeRF-360 [3].

### 3.4. Latent Vector for Tone Disentanglement

In outdoor scenes, it is infeasible to capture training images under controlled conditions, resulting in variations in tone, *i.e.*, white balance and brightness, across shots. The presence of such variations in the training images could lead to inconstancy in the ground truth during training, potentially leading to failed training and artifacts. To mitigate this, we integrate these variations into a latent vector, thereby encouraging the training process to disentangle albedo and tone. As in [25], we pretrain the latent vectors using the approach of Generative Latent Optimization (GLO) [5]. This ensures that each image is assigned to a latent vector  $l$ , which is used as an input to  $M_{albedo}$  and  $\mathbf{c}$  in Eq.(3) is replaced with Eq. (5):

$$\mathbf{c} = M_{albedo}(\mathcal{G}_{albedo}(\mathbf{x}), l, \mathbf{d}). \quad (5)$$

### 3.5. Loss Function

Our SR-TensoRF is optimized based on Mean Squared Error (MSE) loss, L1 loss for density tensor, and Total Variation (TV) loss for all tensors we use. Additionally, a regularization loss is incorporated to constrain shadow values from approaching 1. Please refer to the supplementary materials regarding the weights of the loss terms.

## 4. Experiments

### 4.1. Experimental Details

**Implementation details.** We implement our model based on TensoRF [8]. We utilize shallow MLPs for all three  $M$  components ( $M_{albedo}$ ,  $M_{shadow}$  and  $M_{cubemap}$ ). The numbers of voxels for the density and appearance tensors are the same as in the original TensoRF. The number of elements for the cubemap tensor is fixed at  $128 \times 128$  for a single facet. For optimization, we use the Adam optimizer [20] with a learning rate of 0.01 and a decay of 0.2 for tensor parameters, while the MLPs employed an initial learning rate of 0.001. To prevent overfitting of the cubemap during training, we withhold from integrating the cubemap for the initial 5,000 iterations.

**Datasets and metrics.** We evaluate our approach on NeRF-OSR dataset [32], which contains eight outdoor scene images in total. We show our results for three of them, as in [32]. For the evaluation metrics, we adopt PSNR, Structural Similarity index (SSIM) [44], Mean Absolute Error (MAE), and Mean Squared Error (MSE) of entire test images. For quantitative evaluation, since NeRF-OSR dataset does not include masks, we generate them using the segmentation model SegFormer [46] and exclude the sky

Method	PSNR $\uparrow$	SSIM $\uparrow$	MSE $\downarrow$	MAE $\downarrow$
Site 1				
Yu <i>et al.</i> [51]	17.87	0.378	<u>0.017</u>	0.097
Philip <i>et al.</i> [29]	16.63	0.367	0.023	0.113
NeRF-OSR [32]	<b>18.72</b>	0.468	<b>0.014</b>	<b>0.090</b>
NeRF-OSR $^\dagger$ [32]	<u>18.32</u>	<b>0.623</b>	0.019	<u>0.091</u>
NeRF-OSR $^\ddagger$ [32]	13.33	0.524	0.059	0.177
Ours	17.30	<u>0.542</u>	0.021	0.096
Site 2				
Yu <i>et al.</i> [51]	15.28	0.385	0.032	0.138
Philip <i>et al.</i> [29]	12.34	0.272	0.065	0.2
NeRF-OSR [32]	<u>15.43</u>	<u>0.517</u>	<b>0.029</b>	<u>0.133</u>
Ours	<b>15.63</b>	<b>0.632</b>	<u>0.030</u>	<b>0.111</b>
Site 3				
Yu <i>et al.</i> [51]	15.17	0.376	0.033	0.133
Philip <i>et al.</i> [29]	12.28	0.319	0.062	0.179
NeRF-OSR [32]	<u>16.65</u>	<u>0.501</u>	<b>0.024</b>	<u>0.114</u>
Ours	<b>16.74</b>	<b>0.653</b>	<b>0.024</b>	<b>0.093</b>

Table 1. **Quantitative Results.** Our SR-TensoRF is competitive with other methods.  $\dagger$  use the same 104 training images as ours for training, and  $\ddagger$  use the test set environment map for evaluation. The results of NeRF-OSR are taken from the original paper, while those of  $\dagger$  and  $\ddagger$  are computed following the methodology in the NeRF-OSR paper. Note that the reported metrics may exhibit slight variations due to the absence of publicly available masks.

and tree regions from them. Also, we follow the metric computation as in NeRF-OSR. Regarding the latent vectors, as in [25], we optimize them on the left half of each image in the test set and compute the metrics both on the right half (Tab. 2) and the whole image (Tab. 1).

**Baselines.** We compare our approach with NeRF-OSR [32], the state-of-the-art model and the only existing model to handle novel view synthesis and relighting simultaneously. We also compare with the relighting models [29, 51] that do not perform novel view synthesis. These values are from the NeRF-OSR [32] paper. Note that, in the case of Site 1, only 104 out of 160 training images have information on capture time, thus we trained our model with those 104 images. For a fair comparison, we also trained NeRF-OSR with this image set.

### 4.2. Analysis of SR-TensoRF

**Benefit of sun direction input.** Unlike methods that utilize environment maps for relighting, our approach eliminates the need for preprocessing steps like environment map creation and, if necessary, its conversion to spherical harmonics (as shown in Fig. 1). Additionally, even the captured coordinates, time, and date are readily available metadata in mobile environments, incurring minimal associated costs.



(a) Site 1



(b) Site 2



(c) Site 3

Figure 4. **Qualitative Results.** We visualize novel view synthesis results from four different views. Regarding the lighting of the test set images, our method aligns the time and utilizes a latent vector extracted from the left half of the images. In contrast, NeRF-OSR utilizes default lighting and only conducts novel view synthesis without applying relighting to the test set images. Hence, in these results, the color of the sky remains consistent across all views. Additionally, the shadows cast by buildings on the ground or themselves are much more distinct in our results, highlighting the enhanced clarity in our rendering.

Method	PSNR $\uparrow$	SSIM $\uparrow$	MSE $\downarrow$	MAE $\downarrow$	Time
Site 1					
NeRF-OSR [32]	18.30	0.627	0.018	0.089	32h
Ours	17.60	0.557	0.020	0.092	1h 11m
Site 2					
NeRF-OSR [32]	16.16	0.691	0.029	0.107	61h
Ours	15.83	0.641	0.032	0.114	3h 41m
Site 3					
NeRF-OSR [32]	16.49	0.707	0.025	0.095	58h
Ours	17.27	0.667	0.022	0.087	2h 46m

Table 2. **Results on Half Images.** Evaluation on the right half, which is not used for the latent vector optimization, yields similar quantitative results.



Figure 5. **Cubemap Ablation.** SR-TensoRF extends its rendering capabilities to the sky. As we progress from (a) a model without cubemap and ray contraction to (b) a model with only ray contraction, and finally to (c) a model with both, the training efficiency also increases, masking out unnecessary voxels.

**Benefit of TensorRF.** As shown in Tab. 2, SR-TensoRF achieves notably faster training than NeRF-OSR [32] thanks to our foundation on the TensorRF [8] framework. This extends to the inference process as well. While this relative time efficiency is not directly transferrable to unbounded scenes, we successfully leverage the benefits of TensorRF with the assistance of cubemaps.

**Benefit of cubemap.** As shown in Fig. 5a, without considering unbounded scenes, there are problems: artifacts appear in a place where they should not be (sky, void areas, etc.), or holes are formed. While ray contraction in Fig. 5b mitigates this to some extent, the alpha mask of TensorRF still leads to hole formation. However, upon the incorporation of cubemaps as in Fig. 5c, the holes disappear, resulting in a more natural depiction of the sky. Moreover, this enables better utilization of the alpha mask, consequently reducing training time.

### 4.3. Comparison with Other Methods

Our quantitative results, as shown in Tab. 1, demonstrate that our model generally outperforms previous studies on Site 2 and 3 across various metrics. However, our model’s performance is relatively lower on Site 1, which can be attributed to the prevalence of cloudy training images in this particular scene, hindering proper learning of the sun effect. The results of NeRF-OSR [32] are obtained by applying

a default environment map to all scenes during evaluation. Since NeRF-OSR dataset includes spherical harmonics extracted from environment map images of the test set only for Site 1, we applied this environment map lighting to Site 1 as well. Despite the higher accuracy of the test set environment map compared to the default one for test set images, the results show a decline, highlighting NeRF-OSR’s reliance on lighting for scene reconstruction. Our model, on the other hand, achieves competitive results with NeRF-OSR, even without depending on test set environment maps, and effectively applies lighting to the test set.

In Fig. 4, we show our qualitative results. The rendered images of NeRF-OSR [32] display uniform colors across all scenes and views in the test set due to the use of consistent lighting. The outcomes obtained by individually applying environment maps corresponding to each ground truth can be seen in Fig. 6 (only for Site 1). Fig. 4b demonstrates that building shadows, which are not well-reconstructed in NeRF-OSR, are successfully generated in our model. Furthermore, in Fig. 4c, it is evident that the buildings are rendered with varying brightness based on the direction of sunlight, whether it is facing the front or rear of the structures.

### 4.4. Relighting Results

Fig. 6 shows the qualitative results of relighting. Every first row showcases the results of applying the test set environment map on NeRF-OSR [32] (NeRF-OSR<sup>‡</sup> of Tab. 1). It is evident that the quality is notably compromised compared to the results where the same default environment map was applied to all scenes, as observed in Fig. 4a. The second row presents the results of our SR-TensoRF, with different sun directions and latent vectors applied to the same viewing direction as in the first row. In this case, the latent vectors were randomly selected from the left-half images of the test set. Noticeable variations in lighting, distinct from the reconstructions in Fig. 4a, can be observed. The third row depicts the shadows utilized in the rendering process of the second row.

Fig. 7 shows the shadows rendered for relighting. Given the same viewing direction, it is evident that SR-TensoRF successfully generates shadows that vary over time. For instance, in the case of the buildings in Site 1, which face east, the shadows are appropriately generated as the sun transitions from the southern direction in Fig. 7a, moving towards the southwest in Fig. 7b, and further tilting westward in Fig. 7c. Such results are achievable by directly utilizing the sun direction. Further relighting results for various scenes can be found in the supplementary material.

### 4.5. Ablation Study

To validate the effectiveness of our SR-TensoRF, we conduct an ablation study as in Tab. 3. When comparing the results of experiments where shadow, cubemap, and latent




Figure 6. **Relight Rendering.** SR-TensoRF is capable of producing plausible relit images without using environment maps. Even under new lighting conditions, our approach maintains the reconstruction quality without degradation.

Method	PSNR $\uparrow$	SSIM $\uparrow$	MSE $\downarrow$	MAE $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	MSE $\downarrow$	MAE $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	MSE $\downarrow$	MAE $\downarrow$
	Site 1				Site 2				Site 3			
TensoRF [8]	15.32	0.511	0.034	0.123	13.43	0.583	0.050	0.145	14.35	0.626	0.039	0.122
Ours	<b>17.30</b>	<u>0.542</u>	<b>0.021</b>	<b>0.096</b>	<b>15.63</b>	<b>0.632</b>	<b>0.030</b>	<b>0.111</b>	<b>16.74</b>	<b>0.653</b>	<b>0.024</b>	<b>0.093</b>
(-) shadow	<u>17.21</u>	<b>0.543</b>	<b>0.021</b>	<u>0.097</u>	<u>15.46</u>	<u>0.631</u>	<u>0.031</u>	<u>0.114</u>	16.48	0.651	0.026	0.097
(-) cubemap	16.66	0.541	0.027	0.105	<u>15.46</u>	0.623	0.032	<u>0.114</u>	<u>16.66</u>	<b>0.653</b>	<b>0.024</b>	<b>0.093</b>
(-) latent vector	15.34	0.525	0.035	0.126	14.92	0.623	0.041	0.126	15.83	0.643	0.030	0.106

Table 3. **Ablation.** SR-TensoRF demonstrates a performance drop when individually removing each component.



Figure 7. **Relight by time.** Results when keeping other factors constant and only changing the sun direction. SR-TensoRF can generate varying forms of shadows that change over time.

vector are individually omitted, our full SR-TensoRF, which integrates all the components, exhibits the highest performance. The most substantial decrease in performance occurs when the latent vector is removed, primarily due to the significant impact of tone on quantitative metrics. This is particularly evident for Site 1, where many training images exhibit cloudy weather, accordingly the effects of sunlight are minimal and only the tones tend to vary. Latent vectors, however, alone cannot enable our time-varying relighting capabilities. In addition, small holes in the sky may not have a significant impact on quantitative metrics, yet they can be visually prominent. Addressing this issue with a cubemap solution results in more visually appealing results.

## 5. Conclusion

We propose a sun-aligned outdoor scene relighting model SR-TensoRF, which consists of appearance, den-

sity, and cubemap tensors. Our model utilizes the sun direction for relighting, bypassing the need for environment maps. Additionally, it employs latent vectors to disentangle various tones of images. Our proposed cubemap tensor efficiently combines the fast training speed of TensorRF while effectively handling unbounded scenes. With our SR-TensoRF, we demonstrate that the sun direction can be a viable alternative to environment maps, particularly for outdoor scene relighting. SR-TensoRF not only simplifies the relighting pipeline but also enables the generation of dynamic shadows over time.

**Future work.** Our research can be extended in various directions. As we tackle unbounded scenes, there is potential to model shadows cast by surrounding objects such as trees that might not have been captured in the training images. Furthermore, relighting methods that account for cloudy weather conditions could be explored. Additionally, our approach, initially designed for daytime scenes, could be extended to nighttime scenarios by adapting the lighting direction, similar to sun direction utilization.

**Acknowledgements.** This work was supported by NRF grant (2021R1A2C3006659) and IITP grant (2022-0-00320) funded by Korean Government.



## References

- [1] Benjamin Attal, Eliot Laidlaw, Aaron Gokaslan, Changil Kim, Christian Richardt, James Tompkin, and Matthew O’Toole. Törf: Time-of-flight radiance fields for dynamic scene view synthesis. *Advances in neural information processing systems*, 34:26289–26301, 2021. 1, 2
- [2] Jonathan T Barron, Ben Mildenhall, Matthew Tancik, Peter Hedman, Ricardo Martin-Brualla, and Pratul P Srinivasan. Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5855–5864, 2021. 2
- [3] Jonathan T Barron, Ben Mildenhall, Dor Verbin, Pratul P Srinivasan, and Peter Hedman. Mip-nerf 360: Unbounded anti-aliased neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5470–5479, 2022. 1, 2, 5
- [4] Jonathan T Barron, Ben Mildenhall, Dor Verbin, Pratul P Srinivasan, and Peter Hedman. Zip-nerf: Anti-aliased grid-based neural radiance fields. *arXiv preprint arXiv:2304.06706*, 2023. 2
- [5] Piotr Bojanowski, Armand Joulin, David Lopez-Pas, and Arthur Szlam. Optimizing the latent space of generative networks. In *International Conference on Machine Learning*, pages 600–609. PMLR, 2018. 5
- [6] Mark Boss, Raphael Braun, Varun Jampani, Jonathan T Barron, Ce Liu, and Hendrik Lensch. Nerf: Neural reflectance decomposition from image collections. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12684–12694, 2021. 2
- [7] Alexandra Carlson, Ram Vasudevan, and Matthew Johnson-Roberson. Shadow transfer: Single image relighting for urban road scenes. *arXiv preprint arXiv:1909.10363*, 2019. 3
- [8] Anpei Chen, Zexiang Xu, Andreas Geiger, Jingyi Yu, and Hao Su. Tensorf: Tensorial radiance fields. In *European Conference on Computer Vision*, pages 333–350. Springer, 2022. 1, 2, 3, 5, 7, 8
- [9] Zhiqin Chen, Thomas Funkhouser, Peter Hedman, and Andrea Tagliasacchi. Mobilenerf: Exploiting the polygon rasterization pipeline for efficient neural field rendering on mobile architectures. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16569–16578, 2023. 1, 2
- [10] Sylvain Duchêne, Clement Riant, Gaurav Chaurasia, Jorge Lopez-Moreno, Pierre-Yves Laffont, Stefan Popov, Adrien Bousseau, and George Drettakis. Multi-view intrinsic images of outdoors scenes with an application to relighting. *ACM Transactions on Graphics*, page 16, 2015. 2
- [11] Sara Fridovich-Keil, Alex Yu, Matthew Tancik, Qinhong Chen, Benjamin Recht, and Angjoo Kanazawa. Plenoxels: Radiance fields without neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5501–5510, 2022. 1, 2
- [12] Taku Fujitomi, Ken Sakurada, Ryuhei Hamaguchi, Hidehiko Shishido, Masaki Onishi, and Yoshinari Kameda. Lb-nerf: light bending neural radiance fields for transparent medium. In *2022 IEEE International Conference on Image Processing (ICIP)*, pages 2142–2146. IEEE, 2022. 1
- [13] Yuan-Chen Guo, Di Kang, Linchao Bao, Yu He, and Song-Hai Zhang. Nerfren: Neural radiance fields with reflections. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18409–18418, 2022. 1
- [14] Tom Haber, Christian Fuchs, Philippe Bekaer, Hans-Peter Seidel, Michael Goesele, and Hendrik PA Lensch. Relighting objects from image collections. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 627–634. IEEE, 2009. 2
- [15] Ayaan Haque, Matthew Tancik, Alexei A Efros, Aleksander Holynski, and Angjoo Kanazawa. Instruct-nerf2nerf: Editing 3d scenes with instructions. *arXiv preprint arXiv:2303.12789*, 2023. 1, 2
- [16] Yi-Hua Huang, Yue He, Yu-Jie Yuan, Yu-Kun Lai, and Lin Gao. Stylizednerf: consistent 3d scene stylization as stylized nerf via 2d-3d mutual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18342–18352, 2022. 2
- [17] Brian KS Isaac-Medina, Chris G Willcocks, and Toby P Breckon. Exact-nerf: An exploration of a precise volumetric parameterization for neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 66–75, 2023. 2
- [18] Haian Jin, Isabella Liu, Peijia Xu, Xiaoshuai Zhang, Songfang Han, Sai Bi, Xiaowei Zhou, Zexiang Xu, and Hao Su. Tensorf: Tensorial inverse rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 165–174, 2023. 2
- [19] Natasha Khogade, Tomas Simon, Alexei Efros, and Yaser Sheikh. 3d object manipulation in a single photograph using stock 3d models. *ACM Transactions on graphics (TOG)*, 33(4):1–12, 2014. 2
- [20] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 5
- [21] Zhengqi Li, Simon Niklaus, Noah Snavely, and Oliver Wang. Neural scene flow fields for space-time view synthesis of dynamic scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6498–6508, 2021. 1, 2
- [22] Jingwang Ling, Zhibo Wang, and Feng Xu. Shadowneus: Neural sdf reconstruction by shadow ray supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 175–185, 2023. 2
- [23] Andrew Liu, Shiry Ginossar, Tinghui Zhou, Alexei A Efros, and Noah Snavely. Learning to factorize and relight a city. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part IV 16*, pages 544–561. Springer, 2020. 3
- [24] Linjie Lyu, Ayush Tewari, Thomas Leimkühler, Marc Habermann, and Christian Theobalt. Neural radiance transfer fields for relightable novel-view synthesis with global illumination. In *European Conference on Computer Vision*, pages 153–169. Springer, 2022. 2

- [25] Ricardo Martin-Brualla, Noha Radwan, Mehdi SM Sajjadi, Jonathan T Barron, Alexey Dosovitskiy, and Daniel Duckworth. Nerf in the wild: Neural radiance fields for unconstrained photo collections. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7210–7219, 2021. 5
- [26] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *European Conference on Computer Vision*, pages 405–421, 2020. 1, 2, 3
- [27] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM Transactions on Graphics (ToG)*, 41(4):1–15, 2022. 1, 2
- [28] Michael Niemeyer, Jonathan T Barron, Ben Mildenhall, Mehdi SM Sajjadi, Andreas Geiger, and Noha Radwan. Regnerf: Regularizing neural radiance fields for view synthesis from sparse inputs. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5480–5490, 2022. 2
- [29] Julien Philip, Michaël Gharbi, Tinghui Zhou, Alexei A Efros, and George Drettakis. Multi-view relighting using a geometry-aware network. *ACM Trans. Graph.*, 38(4):78–1, 2019. 3, 5
- [30] Albert Pumarola, Enric Corona, Gerard Pons-Moll, and Francesc Moreno-Noguer. D-nerf: Neural radiance fields for dynamic scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10318–10327, 2021. 1, 2
- [31] Brandon Craig Rhodes. Pyephem astronomy library, 1998. 4
- [32] Viktor Rudnev, Mohamed Elgharib, William Smith, Lingjie Liu, Vladislav Golyanik, and Christian Theobalt. Nerf for outdoor scene relighting. In *European Conference on Computer Vision*, pages 615–631. Springer, 2022. 1, 2, 3, 5, 6, 7, 8
- [33] Seunghyeon Seo, Yeonjin Chang, and Nojun Kwak. Flipnerf: Flipped reflection rays for few-shot novel view synthesis. *arXiv preprint arXiv:2306.17723*, 2023. 2
- [34] Seunghyeon Seo, Donghoon Han, Yeonjin Chang, and Nojun Kwak. Mixnerf: Modeling a ray with mixture density for novel view synthesis from sparse inputs. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20659–20668, 2023. 2
- [35] Rakshith R Shetty, Mario Fritz, and Bernt Schiele. Adversarial scene editing: Automatic object removal from weak supervision. *Advances in Neural Information Processing Systems*, 31, 2018. 2
- [36] Pratul P Srinivasan, Boyang Deng, Xiuming Zhang, Matthew Tancik, Ben Mildenhall, and Jonathan T Barron. Nerv: Neural reflectance and visibility fields for relighting and view synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7495–7504, 2021. 1, 2
- [37] Matthew Tancik, Vincent Casser, Xinchun Yan, Sabeek Pradhan, Ben Mildenhall, Pratul P Srinivasan, Jonathan T Barron, and Henrik Kretzschmar. Block-nerf: Scalable large scene neural view synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8248–8258, 2022. 1, 2
- [38] Zhe Jun Tang, Tat-Jen Cham, and Haiyu Zhao. Able-nerf: Attention-based rendering with learnable embeddings for neural radiance field. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16559–16568, 2023. 1
- [39] Marco Toschi, Riccardo De Matteo, Riccardo Spezialetti, Daniele De Gregorio, Luigi Di Stefano, and Samuele Salti. Relight my nerf: A dataset for novel view synthesis and relighting of real world objects. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20762–20772, 2023. 1, 2
- [40] Edgar Tretschk, Ayush Tewari, Vladislav Golyanik, Michael Zollhöfer, Christoph Lassner, and Christian Theobalt. Non-rigid neural radiance fields: Reconstruction and novel view synthesis of a dynamic scene from monocular video. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12959–12970, 2021. 1, 2
- [41] Alejandro Troccoli and Peter K Allen. Relighting acquired models of outdoor scenes. In *Fifth International Conference on 3-D Digital Imaging and Modeling (3DIM'05)*, pages 245–252. IEEE, 2005. 2
- [42] Haithem Turki, Deva Ramanan, and Mahadev Satyanarayanan. Mega-nerf: Scalable construction of large-scale nerfs for virtual fly-throughs. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12922–12931, 2022. 1, 2
- [43] Dor Verbin, Peter Hedman, Ben Mildenhall, Todd Zickler, Jonathan T Barron, and Pratul P Srinivasan. Ref-nerf: Structured view-dependent appearance for neural radiance fields. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5481–5490. IEEE, 2022. 1
- [44] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004. 5
- [45] Qiling Wu, Jianchao Tan, and Kun Xu. Palettenerf: Palette-based color editing for nerfs. *arXiv preprint arXiv:2212.12871*, 2022. 2
- [46] Enze Xie, Wenhai Wang, Zhiding Yu, Anima Anandkumar, Jose M Alvarez, and Ping Luo. Segformer: Simple and efficient design for semantic segmentation with transformers. *Advances in Neural Information Processing Systems*, 34:12077–12090, 2021. 5
- [47] Zhiwen Yan, Chen Li, and Gim Hee Lee. Nerf-ds: Neural radiance fields for dynamic specular objects. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8285–8295, 2023. 1
- [48] Bangbang Yang, Yinda Zhang, Yinghao Xu, Yijin Li, Han Zhou, Hujun Bao, Guofeng Zhang, and Zhaopeng Cui. Learning object-compositional neural radiance field for editable scene rendering. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 13779–13788, 2021. 2

- [49] Alex Yu, Ruilong Li, Matthew Tancik, Hao Li, Ren Ng, and Angjoo Kanazawa. Plenotrees for real-time rendering of neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5752–5761, 2021. [1](#), [2](#)
- [50] Alex Yu, Vickie Ye, Matthew Tancik, and Angjoo Kanazawa. pixelnerf: Neural radiance fields from one or few images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4578–4587, 2021. [2](#)
- [51] Ye Yu, Abhimitra Meka, Mohamed Elgharib, Hans-Peter Seidel, Christian Theobalt, and William AP Smith. Self-supervised outdoor scene relighting. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXII 16*, pages 84–101. Springer, 2020. [3](#), [5](#)
- [52] Yu-Jie Yuan, Yang-Tian Sun, Yu-Kun Lai, Yuewen Ma, Rongfei Jia, and Lin Gao. Nerf-editing: geometry editing of neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18353–18364, 2022. [1](#)
- [53] Kai Zhang, Nick Kolkin, Sai Bi, Fujun Luan, Zexiang Xu, Eli Shechtman, and Noah Snavely. Arf: Artistic radiance fields. In *European Conference on Computer Vision*, pages 717–733. Springer, 2022. [2](#)
- [54] Kai Zhang, Gernot Riegler, Noah Snavely, and Vladlen Koltun. Nerf++: Analyzing and improving neural radiance fields. *arXiv preprint arXiv:2010.07492*, 2020. [4](#)